CLAIMS

What is claimed is:

1. A method for recognizing speech, the method comprising:

recognizing a sequence of words;

processing the sequence of words using word agglomeration; and

classifying the processed sequence of words as a predetermined command.

2. The method of claim 1, further comprising performing an action corresponding to the predetermined command.

3. The method of claim 1, wherein processing comprises replacing the sequence of words with an associated word *n*-tuple sequence.

4. The method of claim 1, wherein the associated word *n*-tuple sequence is a sequence of all strings of *n* consecutive words present in the sequence of words.

5. The method of claim 3, wherein classifying comprises semantically inferring the predetermined command from the associated word *n*-tuple sequence.

6. The method of claim 1, wherein classifying comprises semantically inferring the predetermined command from the processed sequence of words.

7. The method of claim 6, wherein semantically inferring comprises determining a correlation between the processed sequence of words and at least one semantic anchor.

8. The method of claim 7, wherein the correlation is a distance between a vector corresponding to the processed sequence of words and a vector corresponding to the at least one semantic anchor.

9. The method of claim 8, wherein semantically inferring further comprises selecting the predetermined command from the semantic anchor vector having the shortest distance.

10. The method of claim 9, wherein the semantic anchor represents a one of a plurality of predetermined commands.

11. The method of claim 7, wherein the at least one semantic anchor is derived from a training data.

12. The method of claim 6, wherein semantically inferring the predetermined command depends on the order of the words in the processed sequence of words.

13. The method of claim 1, wherein the classifying comprises:

generating a vector representation of the processed sequence of words; and

comparing the vector representation to a plurality of semantic anchors, wherein each of the plurality of semantic anchors corresponds to one of a plurality of voice commands.

14. The method of claim 13, wherein the classifying further comprises:

selecting a semantic anchor of the plurality of semantic anchors that is most similar to the vector representation; and

classifying the processed sequence of words as the command that corresponds to the selected semantic anchor.

15. The method of claim 14, wherein the selecting comprises:

for each of the plurality of semantic anchors, identifying the similarity between the vector representation and the semantic anchor by calculating the cosine of the angle between the product of the vector representation and a diagonal matrix of singular values and the product of the semantic anchor and the diagonal matrix of singular values; and

selecting the semantic anchor of the plurality of semantic anchors that corresponds to the largest cosine value as the semantic anchor that is most similar to the vector representation.

16. The method of claim 13, wherein the vector representation is an indication of how frequently each of a plurality of word *n*-tuples occurs within the processed sequence of words.

17. The method of claim 16, wherein each of the plurality of semantic anchors is an indication of how frequently each of the plurality of word *n*-tuples occurs with respect to the corresponding command.

18. The method of claim 13, wherein each of the plurality of semantic anchors represents a plurality of different ways of speaking the corresponding command.

19. The method of claim 13, wherein each of the plurality of semantic anchors represents a plurality of different commands having the same words, but in a different order.

20. A machine-readable medium having stored thereon a plurality of instructions that, when executed by a processor, cause the processor to recognize a voice command by:

recognizing a sequence of words;

processing the sequence of words using word agglomeration; and

classifying the processed sequence of words as a predetermined command.

21. The machine-readable medium of claim 20, wherein the plurality of instructions further cause the processor to perform an action corresponding to the predetermined command.

22. The machine-readable medium of claim 20, wherein the plurality of instructions for processing comprises:

replacing the sequence of words with an associated word *n*-tuple sequence.

23. The machine-readable medium of claim 22, wherein the associated word *n*-tuple sequence is a sequence of all strings of *n* consecutive words present in the sequence of words.

24. The machine-readable medium of claim 22, wherein the instructions for classifying comprises semantically inferring the predetermined command from the associated word *n*-tuple sequence.

25. The machine-readable medium of claim 22, wherein the instructions for semantically inferring comprises determining a correlation between a semantic representation of the associated word *n*-tuple sequence and at least one semantic anchor.

26. The machine-readable medium of claim 25, wherein the instructions for determining a correlation comprise determining a distance between a vector corresponding to the semantic representation and a vector corresponding to the at least one semantic anchor.

27. The machine-readable medium of claim 26, wherein the instructions for semantically inferring further comprises selecting the predetermined command from the semantic anchor vector having the shortest distance to the a vector corresponding to the semantic representation.

28. The machine-readable medium of claim 25, wherein the at least one semantic anchor represents a one of a plurality of predetermined commands.

29. The machine-readable medium of claim 25, wherein the at least one semantic anchor is derived from a training data.

30. The machine-readable medium of claim 25, wherein the performance of the instructions for semantically inferring the predetermined command depends on the order of the words in the processed sequence of words.

31. An apparatus for recognizing a voice command, the apparatus comprising:

a speech recognizer to recognize a sequence of words received as the voice command;

a processor to process the sequence of words using word agglomeration; and

a semantic classifier, coupled to the processor, to semantically infer from a vector representation of the processed sequence of words which of a plurality of predetermined commands correlate to the voice command.

32. The apparatus of claim 31, further comprising:

an action generator, coupled to the semantic classifier, to use the vector representation to determine an action to be performed.

33. The apparatus of claim 31, wherein the semantic classifier is further to compare the vector representation to a plurality of semantic anchors, wherein each of the plurality of semantic anchors corresponds to a one of the plurality of predetermined commands.

34. The apparatus of claim 33, wherein the semantic classifier is further to identify a semantic anchor of the plurality of semantic anchors that is most similar to the vector representation, and to classify the vector representation as the one of the plurality of predetermined commands that corresponds to the identified semantic anchor.

35. An apparatus for recognizing a voice command, the apparatus comprising:

means for recognizing a sequence of words received as the voice command;

means for processing the sequence of words using word agglomeration; and

means, coupled to the means for processing, for semantically inferring from a vector representation of the processed sequence of words which of a plurality of predetermined commands correlate to the voice command.

36. The apparatus of claim 35, further comprising means, coupled to the means for semantically inferring, for using the vector representation to determine an action to be performed.

37. The apparatus of claim 35, wherein the means for semantically inferring further comprises:

means for comparing the vector representation to a plurality of semantic anchors, wherein each of the plurality of semantic anchors corresponds to one of a plurality of predetermined commands.

38. The apparatus of claim 37, wherein the means for semantically inferring further comprises:

means for identifying a semantic anchor of the plurality of semantic anchors that is most similar to the vector representation; and

means for classifying the vector representation as the one of the plurality of predetermined commands that corresponds to the identified semantic anchor.

39. A method for characterizing language comprising:

grouping an input sequence of words into at least one of a plurality of word *n*-tuples, wherein each word *n*-tuple is a string of *n* consecutive words present in the sequence;

mapping the plurality of word *n*-tuples into a first vector of a vector space, wherein each word *n*-tuple is an element of the first vector; and

calculating a distance between the first vector and a plurality of second vectors in the vector space, each of said plurality of second vectors having been previously mapped from a training sequence of words.

40. The method of claim 39 wherein grouping the input sequence of words into at least one of a plurality of word *n*-tuples is performed by:

grouping *n* consecutive words of the sequence into a first word *n*-tuple starting with the left-most word in the sequence of words;

repeating the grouping of *n* consecutive words of the sequence into a subsequent word *n*-tuple until all of the words in the sequence of words have been mapped, each repetition starting with the next left-most word in the sequence of words.

41. The method of claim 40, wherein grouping starts with the right-most word in the sequence of words, and each repetition starts with the next right-most word in the sequence of words.

42. The method of claim 39 wherein grouping a sequence of words into at least one of a plurality of word n-tuples is performed by word agglomeration.

43. The method of claim 39 wherein mapping the plurality of word *n*-tuples further includes:

tabulating the number of times each word *n*-tuple occurs in the plurality of word n-tuples using a word *n*-tuple by *n*-tuple document matrix, wherein entries from the matrix form the first vector in the vector space.

44. The method of claim 39 wherein mapping the training sequence of words further includes:

tabulating the number of times a word n-tuple grouped from the training sequence of words occurs in a set of N training documents using a word *n*-tuple by *n*-tuple document matrix, wherein entries from the matrix form the plurality of second vectors in the vector space, which have dimension N.

45. The method of claim 39 further comprising classifying the input sequence of words as representing a predetermined command corresponding to a one of the plurality of second vectors closest in distance to the first vector.